

A Bayesian Hierarchical Model for the Analysis of a Longitudinal Dynamic Contrast-Enhanced MRI Cancer Study

Volker J. Schmid[†] Brandon Whitcher* Anwar R. Padhani[‡]
N. Jane Taylor[†] Guang-Zhong Yang[†]

February 2, 2008

* Corresponding Author:

Clinical Imaging Centre

GlaxoSmithKline

Hammersmith Hospital

Imperial College London

Du Cane Road

London W12 0HS

United Kingdom

+44 20 8008 6052; +44 20 8008 6491 (fax)

brandon.j.whitcher@gsk.com

[†] Institute of Biomedical Engineering

Imperial College London

London SW7 2AZ

United Kingdom

[‡] Paul Strickland Scanner Centre

Mt. Vernon Cancer Centre

Northwood HA6 2RN

United Kingdom

Key words.

Running head. Bayesian Hierarchical Model for DCE-MRI in Oncology

Approximate word count. 5000

Abstract

Imaging in clinical oncology trials provides a wealth of information that contributes to the drug development process, especially in early phase studies. This paper focuses on kinetic modeling in DCE-MRI, inspired by mixed-effects models that are frequently used in the analysis of clinical trials. Instead of summarizing each scanning session as a single kinetic parameter – such as median K^{trans} across all voxels in the tumor ROI – we propose to analyze all voxel time courses from all scans and across all subjects simultaneously in a single model. The kinetic parameters from the usual non-linear regression model are decomposed into unique components associated with factors from the longitudinal study; e.g., treatment, patient and voxel effects. A Bayesian hierarchical model provides the framework in order to construct a data model, a parameter model, as well as prior distributions. The posterior distribution of the kinetic parameters is estimated using Markov chain Monte Carlo (MCMC) methods. Hypothesis testing at the study level for an overall treatment effect is straightforward and the patient- and voxel-level parameters capture random effects that provide additional information at various levels of resolution to allow a thorough evaluation of the clinical trial. The proposed method is validated with a breast cancer study, where the subjects were imaged before and after two cycles of chemotherapy, demonstrating the clinical potential of this method to longitudinal oncology studies.

1 Introduction

Assessing the efficacy of cancer treatments using *in vivo* imaging is shifting from qualitative techniques to quantitative imaging methods that characterize biologically relevant properties of tumor tissue. The use of model-free or heuristic measures, such as the initial area under the Gadolinium curve (IAUGC), or fully quantitative measures, such as the kinetic parameters from a compartmental model, are relatively well understood in the analysis of dynamic contrast-enhanced magnetic resonance imaging (DCE-MRI) (1; 2; 3). Analysis of an oncology imaging trial is usually achieved by applying statistical summaries, such as the mean or median, to the parameters of interest derived from tissue regions of interest (ROIs). That is, enhancing (tumor) voxels are identified from the DCE-MRI data for each scan across all subjects and those voxels are represented by a single parameter; e.g., K^{trans} from quantitative analysis and IAUGC_{90} from a heuristic analysis. Hypothesis testing, either parametric or non-parametric, may then be applied to the derived statistics in order to assess the effects of treatment.

Applying statistical summaries to the kinetic parameter maps from DCE-MRI however discards a substantial amount of information contained in the contrast agent concentration time curves (CTCs) at each voxel, essentially abstracting thousands of observations in space and time to a single number per scan per subject. We believe that there is a wealth of potential information by retaining the collection of CTCs across all subjects and scans, acknowledging the fact that not all CTCs are the same and not all patients are the same.

This paper proposes a Bayesian hierarchical model to analyze all tumor CTCs across all patients and scans in a given study simultaneously based on the concept of a mixed-effects model. Mixed-effects models are well established in the statistical community and have found widespread applications in, for example, agriculture, economics, geophysics and the analysis of clinical trials (4; 5). Mixed-effects models extend the concept of traditional linear or non-linear models by combining both fixed effects and random effects in the same model. More generally, mixed-effects models are most often used to describe relationships between the measured response and explanatory variables in data that are grouped according to one or more factors. Fixed effects denote parameters that are associated with an entire population and random effects denote parameters which are associated with random samples from a population. For example, the drug or radiation therapy given in a trial is a fixed effect because there is no randomness associated with it, whereas patients are inherently random because they are sampled from the population of all patients with a given disease. By acknowledging the fact that some parameters are associated with random samples from a population, we are able to generalize the results from a mixed-effects model beyond the collection of subjects used in the model fitting.

Bayesian methods are used in the construction and estimation of the generalized additive model (6; 7) associated with each kinetic parameter in the model of the CTCs. Similarly to mixed-effects models in a maximum likelihood setting the variances associated with the fixed effects are chosen to be constant, but the variance terms associated with the random effects have prior distributions. This leads to a shrinkage estimation of the random effects so that the random effects are pushed towards zero (8). The fixed effect

in the model explain as much variance as possible, whereas the random effects capture variability that cannot be explained by the fixed effects.

Formulation of a Bayesian hierarchical model is typically achieved in three stages: the data model, the parameter model and the prior parameters (9; 10). The data model reflects our knowledge of the CTCs at the voxel level using the class of compartmental models (11; 12; 13) with a standard arterial input function (AIF) taken from the literature (14; 15; 16). The process model describes how model parameters are generated from underlying processes. At this step we decompose the kinetic parameters into treatment, patient and voxel effects. In Bayesian theory all parameters are regarded as random variables with pre-specified (prior) distributions. This includes the parameters for the fixed as well as the random effects in the model. Conjugate distributions with appropriately selected values are used as priors on all parameters in the model in order to limit their range during the estimation procedure. These choices also allow us to implement efficient sampling methods, wherever possible, to reduce the computational burden.

As an illustration of the parameter model let ψ be a kinetic parameter of interest and suppose the imaging study acquires a dynamic MRI measurement sequence at two time points for each subject – before and after treatment. The generalized additive model uses the natural logarithm as the link function between the signal and effects that can be identified from the study design. For example, assume the imaging study included two time points (pre- and post-treatment) so the parameter ψ may be decomposed via

$$\ln \psi = \text{baseline} + \text{treatment} + \text{patient} + (\text{patient} \cdot \text{treatment}) + \text{voxel}. \quad [1]$$

With model fitting considering all effects related to the kinetic parameter ψ , a curve that fits the observed data at a particular voxel (associated with a specific scanning session and patient) can be derived. This is illustrated in Figure 2 where pre- and post-treatment voxels have been selected from three subjects. The solid line in each plot is the fitted curve to the CTC at a single voxel including all effects in the model (Eq. [1]).

One can construct different versions of ψ that involve a subset of terms in Eq. [1]. For example, defining $\hat{\psi}_b = \exp(\text{baseline})$ produces the estimated baseline value of the kinetic parameter across the entire study, while ignoring patient- and voxel-specific information and $\hat{\psi}_t = \exp(\text{baseline} + \text{treatment})$ produces the estimated treatment value of the kinetic parameter across the entire study when patient- and voxel-specific information are ignored. In Figure 2, the dotted line in each plot is the fitted curve corresponding to the posterior median CTC for the whole study pre-treatment (top row) and post-treatment (bottom row).

Relative changes between baseline and treatment are also available by looking at the individual components in Eq. [1]; e.g., $\hat{\psi}_{t*} = \exp(\text{treatment})$ produces the percentage change due to treatment relative to the baseline value $\hat{\psi}_b$. Given the mathematical properties of exponential functions, we know that $\hat{\psi}_b \cdot \hat{\psi}_{t*} = \hat{\psi}_t$, thus relating the relative changes attributed to specific effects in the generalized additive model to absolute values of the kinetic parameter. Similar manipulations may be performed to investigate patient- or voxel-specific effects and their relative changes from baseline. This figure is described in greater detail in the results section.

The goal of model construction and estimation presented here is to quantify the effect of drug treatment on disease – in this case breast cancer – through quantitative summaries of tumor microvasculature using DCE-MRI. The model framework we have adopted provides a unified treatment of imaging information at the study level through simultaneous estimation of parameters at the voxel, patient and treatment level, allowing a thorough interrogation of the results.

2 Bayesian Hierarchical Model

Bayesian methods rely on the specification of prior distributions $p(\boldsymbol{\theta})$ that express our information about the unknown parameters $\boldsymbol{\theta}$ before any measurements are obtained; i.e., our model assumptions. To assess the model parameters after observing the data, the posterior distribution $p(\boldsymbol{\theta} | \mathbf{Y})$ can be computed, where $\boldsymbol{\theta}$ is the vector of all unknown parameters and \mathbf{Y} is the vector of observations. The posterior distribution of the parameter vector $\boldsymbol{\theta}$ is obtained by applying Bayes’ theorem

$$p(\boldsymbol{\theta} | \mathbf{Y}) = \frac{p(\boldsymbol{\theta}) \ell(\mathbf{Y} | \boldsymbol{\theta})}{\int p(\boldsymbol{\theta}^*) \ell(\mathbf{Y} | \boldsymbol{\theta}^*) d\boldsymbol{\theta}^*}, \quad [2]$$

where $\ell(\mathbf{Y} | \boldsymbol{\theta})$ denotes the likelihood function of \mathbf{Y} and $p(\boldsymbol{\theta})$ the product of all *a priori* probability distribution functions. One can think of the posterior as an update to the prior distribution, our beliefs, on $\boldsymbol{\theta}$ after measuring a process – producing a mixture of previous knowledge and experimental data. Bayesian methods are inherently iterative, since the posterior distribution can become our new prior distribution and, be combined with new measurements of the data generating process at a later date, to produce an updated posterior distribution.

The following sections introduce the key components in the Bayesian hierarchical model: the data model, the parameter model and the prior parameters. Each stage of the model development has been tailored to the analysis of a longitudinal cancer treatment study with two time points. Figure 1 provides a schematic overview of the proposed Bayesian Hierarchical Model (BHM). The three model stages are the rows and the columns represent the “resolution” of the parameters. For example, K^{trans} is decomposed into global (study-wide), subject and voxel effects through the BHM where as v_p is simply estimated for each voxel without further decomposition. The measurement error term is independent of the specific parameter model and involves both prior and hyperprior distributions. A standard compartmental model is used to describe the concentration time curves observed at each voxel. A generalized additive model is proposed to decompose the kinetic parameters into factors that are relevant to the design of the longitudinal study. Finally, the prior distributions, including necessary hyperparameters, are specified on all factors of the parameter model. These prior distributions are flat in most cases, reflecting a lack of knowledge concerning the parameter, but also incorporate biological knowledge, such as a transfer rate must be non-negative, or statistical knowledge, for example a variance must be non-negative.

2.1 Data Model

A hierarchical Bayesian framework is used to model the contrast agent concentration time curve (CTC) of all voxels (17). Let $\mathbf{Y} = [Y(t_1), Y(t_2), \dots, Y(t_T)]^T$ denote the CTC associated with a single voxel observed at T time points determined by the image acquisition protocol. The CTC is assumed to follow a standard compartment model (16)

$$C_t(t) = v_p C_p(t) + C_p(t) \otimes K^{\text{trans}} \exp(-t k_{\text{ep}}), \quad [3]$$

where \otimes denotes the convolution operator, K^{trans} represents the transfer rate from plasma to extracellular extravascular space (EES) per minute, k_{ep} the rate constant between EES and blood plasma per minute and v_p the vascular space fraction. The choice of model for the CTC depends on the scientific goals of the study. Replacing Eq. [3] with a more or less complicated model is straightforward in this model-building framework. The observed vector \mathbf{Y} may therefore be thought of as noisy observations of the true contrast agent concentration $C_t(t)$ given by a draw from a multivariate Normal distribution

$$\mathbf{Y} \sim N_T(\mathbf{C}_t, \sigma^2 I_T), \quad [4]$$

where the notation $Y \sim N(\mu, \sigma^2)$ means that the random variable Y is drawn from the Normal distribution with parameters μ and σ^2 .

We assume a common arterial input function (AIF), taken from the literature for all patients in the study, and we follow the work of Tofts and Kermode (12) by using a bi-exponential function

$$C_p(t) = D[a_1 \exp(m_1 t) + a_2 \exp(m_2 t)], \quad [5]$$

where $a_1 = 24.0$ kg/l, $a_2 = 6.20$ kg/l, $m_1 = 3.00$ min⁻¹ and $m_2 = 0.016$ min⁻¹ are inspired by the work of Fritz-Hansen *et al.* (15).

A Bayesian implementation of the compartmental model above was proposed in Schmid *et al.* (18). Since the Bayesian model framework does not depend on any optimization procedure, it will produce valid parameter estimates when estimation via nonlinear regression fails to converge. Samples from the posterior distribution are built up during the model fitting procedure for each parameter. Hence, the posterior distribution may be used to obtain additional information on the accuracy and precision of the estimates. For example, the standard error of the posterior is the estimation error. Statistics of interest may be derived from the posterior distribution (e.g., mean, median, quantiles, etc.) so that not only point estimates but also confidence intervals are readily available.

2.2 Parameter Model

The pharmacokinetic (PK) parameters from the data model are estimated at every tumor voxel across all subjects and scans. We assume *a priori* that the distribution of the random variables K^{trans} and k_{ep} in the tumor are patient-specific and are changed by treatment in a similar way. Therefore a generalized additive model is used where the log-transformed kinetic parameters $\ln(K^{\text{trans}})$ and $\ln(k_{\text{ep}})$ are expressed

as a linear combination of fixed- and random-effects associated with identifiable factors in the study. In addition to mathematical convenience, the log transform is also appealing in this context because individual terms in the additive model may be interpreted as percentage change from baseline. We assume that the distribution of the vascular fraction v_p will not be changed by the treatment, however single v_p values will be changed. Let $i = 1, \dots, I$ denote the scans acquired and let $j = 1, \dots, J$ denote the patients, so that n_{ij} denotes the number of tumor voxels for patient j at scan i , measured at T time points. The transfer rate constants in Eq. [3] are assumed to be non-negative and estimated in log-transformed space. That is, let $\psi_1 = \ln(K^{\text{trans}})$, $\psi_2 = \ln(k_{\text{ep}})$ and $\boldsymbol{\psi} = [\ln(K^{\text{trans}}), \ln(k_{\text{ep}})]^T$.

The factor of interest when measuring a change in the kinetic parameters is the treatment effect, or the difference between scan $i = 1$ and scan $i = 2$ when only pre- and post-treatment images are acquired. We acknowledge the fact that substantial variability exists across patients in the study and between the voxels in each region of interest (ROI) that describes the enhancing region in the acquisition. Hence, the model for $\ln(K^{\text{trans}})$ is given by

$$\psi_{ijk1} = [1 \ x_i] \begin{bmatrix} \alpha_1 \\ \beta_1 \end{bmatrix} + [1 \ x_i] \begin{bmatrix} \gamma_{j1} \\ \delta_{j1} \end{bmatrix} + \epsilon_{ijk1}, \quad \text{for all } i, j, k \quad [6]$$

where

$$x_i = \begin{cases} 1 & \text{scan } i = 2; \\ 0 & \text{otherwise.} \end{cases} \quad [7]$$

The parameter α_1 is the value of $\ln(K^{\text{trans}})$ associated with the baseline scan and β_1 is the treatment effect (since it is only associated with the post-treatment acquisition). These parameters are regarded as fixed effects (the global column of Figure 1), and thus do not vary between patients in the study. In the Bayesian framework, a marginal posterior distribution will be available for each parameter. The parameter γ_{j1} is the effect of patient j on $\ln(K^{\text{trans}})$ and δ_{j1} is the interaction between patient j and treatment. These parameters are random effects since each patient is assumed to be drawn from a larger population of patients suffering from this condition (the subject column of Figure 1). Finally, the parameter ϵ_{ijk1} is the random effect of voxel k in scan i of patient j on $\ln(K^{\text{trans}})$. The voxel effect acknowledges the fact that each voxel in the tumor volume is drawn from a distribution that describes the ideal tumor voxel (the voxel column of Figure 1). The combination of fixed and random effects in a single model is commonly referred to as a mixed-effects model (6).

Using matrix notation, we can combine the generalized additive model across both kinetic parameters, $\ln K^{\text{trans}}$ and $\ln k_{\text{ep}}$, such that

$$\boldsymbol{\psi}_{ijkl} = \mathbf{Z}_i \begin{bmatrix} \boldsymbol{\phi} \\ \boldsymbol{\theta}_j \end{bmatrix} + \boldsymbol{\epsilon}_{ijkl} \quad [8]$$

$$\mathbf{X}_i = \begin{bmatrix} 1 & x_i & 0 & 0 \\ 0 & 0 & 1 & x_i \end{bmatrix}; \quad \mathbf{Z}_i = [\mathbf{X}_i \ \mathbf{X}_i]; \quad \boldsymbol{\phi}_l = \begin{bmatrix} \alpha_1 \\ \beta_1 \\ \alpha_2 \\ \beta_2 \end{bmatrix}; \quad \boldsymbol{\theta}_{jl} = \begin{bmatrix} \gamma_{j1} \\ \delta_{j1} \\ \gamma_{j2} \\ \delta_{j2} \end{bmatrix}; \quad \boldsymbol{\epsilon}_{ijkl} = \begin{bmatrix} \epsilon_{ijk1} \\ \epsilon_{ijk2} \end{bmatrix} \quad [9]$$

The scan-specific covariates in the model are captured in \mathbf{Z}_i , the fixed effects are in ϕ_l , the patient-specific random effects are in θ_{jl} and the voxel-specific random effects are in ϵ_{ijkl} . The model formulation in Eq. [8] can be adapted in order to incorporate a greater number of scans in a longitudinal study.

2.3 Prior Models

In the Bayesian framework prior information with unknown variance is used to model the random effects. We use vector notation to denote the patient-specific random effects such that

$$\boldsymbol{\gamma} = \begin{bmatrix} \gamma_{11} \\ \gamma_{21} \\ \vdots \\ \gamma_{J1} \\ \gamma_{12} \\ \vdots \\ \gamma_{J2} \end{bmatrix} \quad \text{and} \quad \boldsymbol{\delta} = \begin{bmatrix} \delta_{11} \\ \delta_{21} \\ \vdots \\ \delta_{J1} \\ \delta_{12} \\ \vdots \\ \delta_{J2} \end{bmatrix}, \quad [10]$$

where we have dropped the kinetic parameter subscript to simplify the notation. We draw from multivariate Gaussian distributions to characterize the prior distribution of the unknown variances for the patient-specific random effects, i.e.,

$$\boldsymbol{\gamma} \sim \text{N}_{2J}(\mathbf{0}, \text{diag}(\boldsymbol{\tau}_{\gamma}^2)), \quad [11]$$

$$\boldsymbol{\delta} \sim \text{N}_{2J}(\mathbf{0}, \text{diag}(\boldsymbol{\tau}_{\delta}^2)), \quad [12]$$

where $\boldsymbol{\tau}_{\gamma}^2$ and $\boldsymbol{\tau}_{\delta}^2$ are vectors of the same length and indexed as $\boldsymbol{\gamma}$ and $\boldsymbol{\delta}$, respectively. The voxel-specific random-effect vectors are given unique prior distributions by scan, patient and parameter, so that each vector is given by $\boldsymbol{\epsilon}_{ijl} = [\epsilon_{ij1l}, \epsilon_{ij2l}, \dots, \epsilon_{ijn_{ij}l}]^T$ and it is drawn from a multivariate Gaussian distribution via

$$\boldsymbol{\epsilon}_{ijl} \sim \text{N}_{n_{ij}}(\mathbf{0}, \tau_{\epsilon;ijl}^2 I_{n_{ij}}) \quad [13]$$

where n_{ij} is the number of voxels in the region of interest of scan i of patient j , and $\tau_{\epsilon;ijl}^2$ is the unknown variance associated with scan i , patient j and parameter l . Since the variances are unknown parameters, they must have their own prior distributions which are given by

$$\boldsymbol{\tau}_{\gamma}^2 \stackrel{\text{iid}}{\sim} \text{IG}(1, 1), \quad [14]$$

$$\boldsymbol{\tau}_{\delta}^2 \stackrel{\text{iid}}{\sim} \text{IG}(1, 1), \quad [15]$$

$$\tau_{\epsilon}^2 \stackrel{\text{iid}}{\sim} \text{IG}(1, 10^{-5}), \quad [16]$$

where $\text{IG}(a, b)$ denotes the Inverse Gamma distribution (19), allowing only non-negative values. The inverse Gamma distribution is a conjugate prior for the Normal distribution. For the fixed effects we use flat priors; i.e., the prior distribution does not contain any relevant information, such that

$$p(\alpha_l) = p(\beta_l) = \text{constant} \quad \text{for } l = 1, 2. \quad [17]$$

The prior distributions on the coefficients in the generalized additive model are chosen so that as much variance in the data is explained by the fixed effects α and β – as no prior information is used for those parameters. Variability which cannot be explained by the fixed effects will be covered by the random effects γ and δ . For these parameters an additional prior distribution (hyperprior) on the variance of the parameters is defined, which leads to a shrinkage of those effects, that is the parameters are pushed towards zero and therefore do not cover variance explained by the fixed effects. Any additional variance is explained by the voxel effects.

For the vascular space fraction we impose a relatively flat prior

$$v_{p;ijk} \stackrel{\text{iid}}{\sim} \text{B}(1, 19), \quad \text{for all } i, j, k, \quad [18]$$

where $\text{B}(a, b)$ denotes the Beta distribution (20), so that the *a priori* expected value of v_p is 0.05. The Bayesian hierarchical model is complete by specifying a prior distribution for the variance of the observational error in Eq. [4], with one variance parameter per scan per patient,

$$\sigma_{ij}^2 \stackrel{\text{iid}}{\sim} \text{IG}(1, 10^{-2}) \quad \text{for all } i, j. \quad [19]$$

3 Materials and Methods

3.1 Data acquisition

The first twelve patients from a previously reported breast cancer study are included in the analysis (21; 18). Data were provided by the Paul Strickland Scanner Centre (PSSC) at Mount Vernon Hospital, Northwood, UK. Each patient underwent a DCE-MRI study before and after two cycles of chemotherapy (5-fluorouracil, epirubicin and cyclophosphamide). Six of these patients were identified as pathological responders after receiving six cycles of chemotherapy, the others were non-responders.

For the calculation of T_1 values, we used a two-point measurement with calibration curves as described in (22; 23). The T_1 values are computed as ratio of a T_1 -weighted fast low-angle shot (FLASH) image and a proton density weighted (PDw) FLASH image. The imaging parameters of the T_1 -weighted FLASH images were $\text{TR} = 11$ ms, $\text{TE} = 4.7$ ms, $\alpha = 35^\circ$, the parameters of the proton density-weighted image were $\text{TR} = 350$ ms, $\text{TE} = 4.7$ ms, $\alpha = 6^\circ$. Field of view was the same for all scans, $260 \times 260 \times 8$ mm per slice, so voxel dimensions were $1.016 \times 1.016 \times 8$ mm. A scan consists of three sequential slices of 256×256 voxels and one slice placed in the contra lateral breast as control, which we do not use for our analysis. A total of 40 to 50 acquisitions were acquired, with one acquisition each 11.9 seconds. A dose of $D = 0.1$ mmol per kg body weight of Gd-DTPA was injected after the fourth scan using a power injector with 4 ml/s with a 20 ml saline flush also at 4 ml/s. The first four scans, before contrast, were used to compute T_{10} as the average of the T_1 values of these images. Data from this study were acquired in accordance with the recommendation given by (24). Informed consent was obtained from all patients.

Regions of interest (ROIs) were drawn manually by an expert radiologist on a scan-by-scan basis using anatomical images and subtraction images from the dynamic data to define tumor voxels in pre and post treatment scans.

3.2 Parameter Estimation via MCMC

The joint posterior distribution of all parameters was assessed using Markov chain Monte Carlo (MCMC) (25). After a initial burn-in phase of 10,000 iterations, another 100,000 iterations were computed. To ensure independent samples only each 100th sample was used, giving us a total of 1000 samples to describe the posterior distribution. The global parameters ϕ and patient-specific θ_j were drawn *en bloc* in Gaussian Gibbs steps (26), and hyperparameters were drawn in independent Gamma Gibbs steps; technical details can be found in Appendix A.1. Metropolis-Hastings steps with random walk proposals were necessary for the voxel-specific random effects and vascular space fraction. The algorithm was tuned to an acceptance rate of 30-50% (27). Summary statistics were computed from the samples of the posterior distribution to provide point estimates of the parameters. Empirical standard errors, along with sample quantiles, were used to characterize the precision of the parameter estimates.

4 Results

All parameter estimates are derived from the posterior distribution using Bayes theorem. Hence, a sampling distribution for each parameter value has been built up from which we can produce a point estimate via the median of the sample and also credible intervals (Bayesian confidence intervals) by using the quantiles from their sampling distributions.

How the individual parameters from the generalized additive model coalesce to fit the observed contrast agent concentration time curve is illustrated, at the voxel level, in Figure 2. The observed CTCs for two voxels from three subjects, one voxel at baseline and one voxel after treatment, are plotted along with three fitted curves. The best estimate from the Bayesian hierarchical model at a specific voxel is provided by the solid lines in each plot. That is, all parameters from the generalized additive model (Eq. [8]) are used in the parameter model in order to fit the data model. These curves are very similar to, but not exactly the same as, model fits from the standard non-linear regression method used in the quantitative analysis of DCE-MRI data (18). Removing the voxel-specific term from the model produces a fitted curve that is associated with patient and treatment effects, but not the specific voxel, and are plotted as dashed lines in Figure 2. Given the presence of inter-voxel heterogeneity in the tumor ROI, the dashed lines may or may not fit the observed data at a given voxel very well but they do represent the best (in the sense of a posterior median) fit to all voxels in the tumor ROI for a given patient at a single scan time point. Going back one more level in the generalized additive model and removing the patient effect leaves a fitted curve associated with the baseline and post-treatment scans (i.e., two curves that summarize the overall treatment effect) given by the dotted lines. The top row of Figure 2 contains voxels from three subjects before treatment so the dotted lines are identical and represent the best (in the sense of a posterior median) fit to all pre-treatment voxels across all subjects. The bottom row contains voxels from the same subjects after treatment and the dotted line is the best fit to all post-treatment voxels.

Figure 3 shows the posterior distributions of pre-treatment (baseline) K^{trans} and post-treatment

K^{trans} . That is, the posterior samples were transformed via $\exp(\alpha_1)$ and $\exp(\alpha_1 + \beta_1)$, respectively. For ease of comparison between the two posterior distributions a smoothed version of each histogram, known as a kernel density estimate, is displayed (28). The posterior median of K^{trans} at baseline is 0.205, and the posterior median of K^{trans} after treatment is 0.156. Credible intervals for K^{trans} , that cover 95% of the posterior distribution, are $[0.186, 0.234]$ at baseline and $[0.121, 0.198]$ after treatment. A credible interval is a posterior probability interval. That is, the true value of K^{trans} lies in the interval $[0.186, 0.234]$ with posterior probability 0.95 at baseline and in $[0.121, 0.198]$ with posterior probability 0.95 after treatment.

The density estimates in Figure 3 are unimodal and indicate an overall decrease in K^{trans} after treatment. In order to test for a treatment effect on K^{trans} , specifically a reduction in K^{trans} in the second acquisition compared to the first, we construct the hypothesis

$$H_0 : \beta_1 > 0 \quad \text{versus} \quad H_1 : \beta_1 \leq 0, \quad [20]$$

using the treatment effect from the parameter model (Eq. [8]) and calculate the posterior probability of β_1 exceeding zero. From the results of the MCMC simulation, the null hypothesis (Eq. [20]) is rejected with $p = 0.001$.

When introducing the generalized additive model previously the fact that the parameter K^{trans} and the covariates are linked through a logarithmic transform leads to the interpretation of individual covariates in the parameter model as percentage changes from baseline instead of absolute changes. For the treatment effect this translates into a $100\% \cdot |0.7659 - 1| = 23.3\%$ median reduction in K^{trans} from baseline, where the sign determines whether the change is associated with an increase or decrease.

Figure 4 shows the patient-specific posterior distributions for pre-treatment K^{trans} , given by $\exp(\alpha_1 + \gamma_{j1})$ for $j = 1, \dots, 12$, and post-treatment K^{trans} , given by $\exp(\alpha_1 + \beta_1 + \gamma_{j1} + \delta_{j1})$ for $j = 1, \dots, 12$. The clinical responders are grouped in the first two columns of Figure 4 and the clinical non-responders are in the third and fourth columns. The same range for x -axis $[0, 0.45]$ was used in all plots of K^{trans} for comparison. In general the decrease in K^{trans} observed in the clinical responders is greater than the clinical non-responders, but this is not absolute. For example, patient 12 shows only a small decrease in K^{trans} post-treatment and patient 6 shows an increase in K^{trans} after treatment, but both are clinical responders after additional chemotherapy. The interpretation of the treatment effect as a percentage change from baseline helps to quantify the results in Figure 4. The median percentage change in K^{trans} for subject j is obtained via $100\% \cdot |\exp(\hat{\beta}_1 + \hat{\delta}_{j1}) - 1|$, where the sign determines whether an increase or decrease occurred. For example, patient 1 (pathological responder) experienced a $100\% \cdot |0.7684 - 1| = 23.2\%$ median reduction in K^{trans} which is very similar to the overall treatment effect. This is definitely not the norm as patient 9 experienced a $100\% \cdot |0.4285 - 1| = 57.2\%$ median reduction in K^{trans} and patient 6 experienced a $100\% \cdot |1.0817 - 1| = 8.17\%$ median *increase* in K^{trans} , both were pathological responders.

Figure 5 shows the voxel-specific median posterior for pre- and post-treatment K^{trans} . The clinical responders are grouped in the first two columns and the clinical non-responders are in the third and fourth columns (identical to Fig. 4). The range for the x -axis was restricted to $[0, 1]$ in all plots for comparison. Given the number of samples from the posterior distribution across all voxels, the median

value of $\exp(\alpha_1 + \gamma_{j1} + \epsilon_{1jk1})$ for $j = 1, \dots, 12; k = 1, \dots, n_{1j}$ and $\exp(\alpha_1 + \beta_1 + \gamma_{j1} + \delta_{j1} + \epsilon_{2jk1})$ for $j = 1, \dots, 12; k = 1, \dots, n_{2j}$ across the 1000 samples, for each voxel k , was computed to summarize the voxel effect. The resulting histograms for the voxel effect have been summarized by a kernel density estimate. Most voxel-level distributions of median K^{trans} show a substantial change in shape after treatment, although this is more apparent in the responders compared to the non-responders. It is interesting to note the extent of changes in the shape of these distributions between the different subjects. For example, patient 11 is characterized by a tumor with two distinct modes in estimated K^{trans} at baseline and a single mode after treatment. Even more interesting is the fact that the post-treatment distribution of K^{trans} is in between the two modes at baseline. The distributions of median K^{trans} for patient 12 show the reverse effect, albeit much more subtle than patient 11, where the post-treatment distribution of median K^{trans} appears to be bimodal but still spans a similar range of values.

5 Discussions and conclusions

Information is obtained at multiple levels during an imaging study in the clinical trial setting. The main scientific question of interest is usually, was there a treatment effect? This key hypothesis test drives study design by influencing critical experimental design parameters such as power and sample size. However, information at other levels, such as the patient or voxel level, can provide insight into much more subtle features concerning patients, tumors and the treatment effect. Patient variability with application to predicting clinical response and tumor heterogeneity, as measured by voxel-wise properties of the pharmacokinetic model, are just two examples of so-called secondary endpoints.

The Bayesian hierarchical model presented here was developed to test the hypothesis of a treatment effect for an imaging study while acknowledging known sources of uncertainty; e.g., patients and voxels. This is similar to the approach taken in standard analysis methods for clinical trials where fixed and random effects are identified in the model. The specification of fixed and random effects allows the results from the study to be applicable beyond the specific patient population recruited for this specific study.

A standard analysis was performed on the ROIs and the median K^{trans} values have been summarized in Table 1. A non-parametric test (a one-sided Wilcoxon signed rank test) was performed to test that the difference between the median values was greater than zero; i.e., the treatment did not reduce K^{trans} across all subjects. The null hypothesis was rejected at a borderline significance level ($p = 0.055$). Given the small sample size, $N_1 = 6$ responders and $N_2 = 6$ non-responders, this is an impressive result and there is obviously a reasonable difference in K^{trans} between the two groups.

Figure 6 shows the kernel density estimates of K^{trans} for each ROI, before and after treatment, using a voxel-wise non-linear regression analysis. That is, the compartmental model in Eq. [3] was fit to each voxel independently using the Levenberg-Marquardt optimization procedure. The empirical distributions observed for each patient are extremely similar to those obtained in the BHM. This is to be expected given the relatively flat priors that were imposed on the kinetic parameters (18).

While the voxel-wise results from the Bayesian and regression methods are very similar, and thus provide a check on the consistency of the Bayesian model fitting procedure, the advantages of the Bayesian hierarchical model are clear through the coefficients from the generalized additive model (Eq. [8]). The regression analysis can only summarize the study through Table 1, but the BHM allows one to isolate and interrogate specific effects, at the study or patient or voxel level, through the generalized additive model. Examples of such interrogations have been presented here in Figures 3 and 4, but the possibilities for such model summaries are only limited by the construction of the parameter model.

Bayesian models rely on *a priori* beliefs about the model and parameters, expressed as prior distributions. In general, a flat prior provides similar information to a maximum likelihood approach, and hence similar results. However, in the Bayesian hierarchical model proposed here the choice of prior distribution is critical in specifying the model. We used flat priors on the baseline α and on the treatment effect β , and thus the approach is similar to a non-Bayesian or frequentist approach. For the patient specific effects γ and δ we used Gaussian priors with unknown variances; this is also known as shrinkage prior, as it shifts the parameters towards zero. Hence the patient-specific effects only pick up the deviation from baseline and treatment effect. The voxel effect was also given a shrinkage prior with a more informative hyperprior distribution on the variance, hence it only picks up variability after modelling the baseline, treatment and patient-specific effects.

In this paper a generalized additive model was constructed for the kinetic parameters (K^{trans} and k_{ep}) in a compartmental model. This model incorporated two scanning sessions, and all subjects, to assess the effect of treatment. The modeling framework is easily extended to handle additional covariates or scanning sessions. For example, a dose-ranging study design could be incorporated into the additive model where the treatment effect can be expressed as a function of the dose. Additional scans over time would enable the assessment of temporal dependence on treatment and provide information about the reliability of the data by potentially reducing the amount of uncertainty in the parameter estimates.

Another possible extension of this model would be to include the spatial information of adjacent voxels. In the current implementation of the Bayesian hierarchical model all voxels from one region of interest (tumor) were treated as spatially independent. Since voxel borders are arbitrary and do not represent physiological boundaries between different tissue types, it is likely that neighboring voxels share similar perfusion characteristics. This fact has been taken advantage of in the context of Bayesian modeling of individual scans from a DCE-MRI study (18). The inclusion of a neighborhood structure in the modeling process would reduce the uncertainty in estimation and provide more reliable estimates of the kinetic parameters.

6 Acknowledgements

Support for VJ Schmid was financed through a research grant from GlaxoSmithKline.

A Appendix

A.1 Full conditional distributions

In each iteration of the MCMC (Markov chain Monte Carlo) algorithm, a random sample of the marginal posterior distribution for all parameters is drawn. This is performed by drawing from the conditional posterior distribution of one or more parameters given all other parameters and the data. Hence, the full conditional distributions must be computed. The full conditional is denoted by $\theta | \cdot$, where θ is the parameter and \cdot denotes all other parameters and the data. If the full conditional takes the form of a standard distribution, one can sample directly from this distribution; this is known as the Gibbs sampler (25). If the full conditional is not a standard distribution, then a Metropolis-Hastings sampler must be constructed.

In the proposed Bayesian hierarchical model all full conditionals are from standard distributions due to the use of conjugate prior distributions, except for the voxel effect and v_p . Let $\boldsymbol{\xi}_l = (\alpha_l, \beta_l, \boldsymbol{\gamma}_l, \boldsymbol{\delta}_l)$ denote the vector of length $P = I(J+1)$ associated with all parameters in the generalized additive model, except the voxel effect, for a specific kinetic parameter. The full conditional of $\boldsymbol{\xi}_l$ is a multivariate Normal distribution given by

$$\begin{aligned}\boldsymbol{\xi}_l | \cdot &\sim N_P(\mathbf{V}^{-1}\mathbf{m}, \mathbf{V}^{-1}), \\ \mathbf{m} &= [m_1, \dots, m_P]^T, \\ m_p &= \sum_{i=1}^2 \sum_{j=1}^J \left(\tau_{\epsilon;ij} \sum_{k=1}^{n_{ij}} w_{ijp} \psi_{ijkl} \right) \quad \text{for } p = 1, \dots, P, \\ \mathbf{V} &= \mathbf{W}^T \Lambda \mathbf{W} + \text{diag}(0, 0, \tau_{\gamma;1l}, \dots, \tau_{\gamma;Jl}, \tau_{\delta;1l}, \dots, \tau_{\delta;Jl}),\end{aligned}$$

where \mathbf{W} is a $I(J+1) \times P$ matrix indicating which covariate should be included in the parameter model (Eq. [8]) and Λ is a diagonal matrix with elements $n_{ij}\tau_{\epsilon;ij}$. The vector $\boldsymbol{\xi}_l$ is drawn in one block from a multivariate Normal distribution with an efficient block-sampling algorithm (29).

The full conditional distribution of the voxel effect ϵ_{ijkl} is a non-standard distribution. For computational reasons it is more convenient to sample from ψ_{ijkl} rather than from ϵ_{ijkl} . where the full conditional distribution of ψ_{ijkl} is given by

$$p(\psi_{ijkl} | \cdot) \propto \exp \left(-\frac{1}{2} \tau_{\epsilon;ijk} \psi_{ijkl}^2 - \frac{1}{2\sigma_{jk}^2} (Y_{ijkl} - \hat{Y}_{ijkl})^2 \right). \quad [21]$$

Note, \hat{Y}_{ijkl} is the estimated contrast agent concentration curve given by the estimated model parameters in $\hat{\psi}_{ijkl}$. Draws from this distribution are obtained using a Metropolis-Hastings step.

The full conditionals of all variance parameters are inverse Gamma distributions, which are given by

$$\tau_{\gamma}^2 | \cdot \stackrel{\text{iid}}{\sim} \text{IG}(1.5, 1 + \gamma_{jl}^2), \quad [22]$$

$$\tau_{\delta}^2 | \cdot \stackrel{\text{iid}}{\sim} \text{IG}(1.5, 1 + \delta_{jl}^2), \quad [23]$$

$$\tau_{\epsilon}^2 | \cdot \stackrel{\text{iid}}{\sim} \text{IG} \left\{ 1 + \frac{1}{2} \sum_{i=1}^I \sum_{j=1}^J n_{ij}, 10^{-5} + \frac{1}{2} \sum_{i=1}^I \sum_{j=1}^J \sum_{k=1}^{n_{ij}} \left(\mathbf{z}_i \begin{bmatrix} \phi \\ \boldsymbol{\theta}_j \end{bmatrix} - \psi_{ijkl} \right)^2 \right\}. \quad [24]$$

Hence, the variance parameters can be drawn independently.

References

- [1] Padhani AR, Yarnold J, Regan J, Husband JE. Dynamic MRI of breast hardness following radiation treatment. *Journal of Magnetic Resonance Imaging* 2003;17:427–434.
- [2] Collins DJ, Padhani AR. Dynamic magnetic resonance imaging of tumor perfusion. *IEEE Engineering in Biology and Medicine Magazine* 2004;65–83.
- [3] Tofts PS, Brix G, Buckley DL, Evelhoch JL, Henderson E, Knopp MV, Larsson HBW, Lee TY, Mayr NA, Parker GJM, Port RE, Taylor J, Weiskoff R. Estimating kinetic parameters from dynamic contrast-enhanced T₁-weighted MRI of a diffusable tracer: Standardized quantities and symbols. *Journal of Magnetic Resonance Imaging* 1999;10:223–232.
- [4] Pinheiro JC, Bates DM. *Mixed-Effects Models in S and S-PLUS*. New York: Springer-Verlag, 2000.
- [5] Brown H, Prescott R. *Applied Mixed Models in Medicine*. Chichester, UK: John Wiley & Sons, 1999.
- [6] Fahrmeir L, Lang S. Bayesian inference for generalized additive mixed models based on Markov random field priors. *Applied Statistics* 2001;50:201–220.
- [7] Fahrmeir L, Tutz G. *Multivariate Statistical Modelling Based on Generalized Linear Models*. New York: Springer, 2nd edition, 2001.
- [8] Tibshirani R. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society Series B* 1996;58:267–288.
- [9] Gelman A, Carlin JB, Stern HS, Rubin DB. *Bayesian Data Analysis*. Boca Raton, Florida: Chapman & Hall/CRC, 2nd edition, 2003.
- [10] Wikle CK. Hierarchical models in environmental science. *International Statistical Review* 2002; 71:181–199.
- [11] Kety S. Blood–tissue exchange methods. Theory of blood-tissue exchange and its application to measurement of blood flow. *Methods in Medical Research* 1960;8:223–227.
- [12] Tofts PS, Kermode AG. Measurement of the blood-brain barrier permeability and leakage space using dynamic MR imaging. 1. Fundamental concepts. *Magnetic Resonance in Medicine* 1984; 17(2):357–367.
- [13] Larsson HB, Tofts PS. Measurement of the blood-brain barrier permeability and leakage space using dynamic Gd-DTPA scanning—a comparison of methods. *Magnetic Resonance in Medicine* 1992; 24(1):174–176.
- [14] Weinmann HJ, Laniado M, Mutzel W. Pharmacokinetics of Gd-DTPA/dimeglumine after intravenous injection into healthy volunteers. *Physiological Chemistry and Physics and Medical NMR* 1984;16:167–172.

- [15] Fritz-Hansen T, Rostrup E, Larsson HBW, S ndergaard L, Ring P, Henriksen O. Measurement of the arterial concentration of Gd-DTPA using MRI: A step toward quantitative perfusion imaging. *Magnetic Resonance in Medicine* 1996;36:225–231.
- [16] Buckley DL, Parker GJM. Measuring contrast agent concentration in T_1 -weighted dynamic contrast-enhanced MRI. In A Jackson, DL Buckley, GJM Parker, eds., *Dynamic Contrast-Enhanced Magnetic Resonance Imaging in Oncology*. Berlin: Springer, 2005; 69–80.
- [17] Gelman A, Stern H, Carlin J, Rubin D. *Bayesian Data Analysis*. Chapman & Hall/CRC Press, London/Boca Raton, 2003.
- [18] Schmid VJ, Whitcher B, Padhani AR, Taylor NJ, Yang GZ. Bayesian methods for pharmacokinetic models in dynamic contrast-enhanced magnetic resonance imaging. *IEEE Transactions on Medical Imaging* 2006;25(12):1627–1636.
- [19] Johnson NL, Kotz S, Balakrishnan N. *Continuous Univariate Distributions*, volume 1. New York: John Wiley & Sons, 2nd edition, 1994.
- [20] Johnson NL, Kotz S, Balakrishnan N. *Continuous Univariate Distributions*, volume 2. New York: John Wiley & Sons, 1995.
- [21] Ah-See MLW, Makris A, Taylor NJ, Burcombe RJ, Harrison M, Stirling JJ, Richman PI, Leach MO, Padhani AR. Does vascular imaging with MRI predict response to neoadjuvant chemotherapy in primary breast cancer? *Journal of Clinical Oncology (Meeting Abstracts)* 2004;22(14S):582.
- [22] Parker G, Suckling J, Tanner S, Padhani A, Revell P, Husband J, Leach M. Probing tumor microvasculature by measurement, analysis and display of contrast agent uptake kinetics. *Journal of Magnetic Resonance Imaging* 1997;7:564–574.
- [23] d’Arcy J, Collins D, Padhani A, Walker-Samuel S, Suckling J, Leach M. Informatics in radiology (infoRAD): Magnetic resonance imaging workbench: analysis and visualization of dynamic contrast-enhanced MR imaging data. *Radiographics* 2006;26(2):621–632.
- [24] Leach MO, Brindle KM, Evelhoch JL, Griffiths JR, Horsman MR, Jackson A, Jayson GC, Judson IR, Knopp MV, Maxwell RJ, McIntyre D, Padhani AR, Price P, Rathbone R, Rustin GJ, Tofts PS, Tozer GM, Vennart W, Waterton JC, Williams SR, Workman P. The assessment of antiangiogenic and antivascular therapies in early-stage clinical trials using magnetic resonance imaging: issues and recommendations. *British Journal of Cancer* 2005;92:1599–1610.
- [25] Gilks WR, Richardson S, Spiegelhalter DJ, eds. *Markov Chain Monte Carlo in Practice*. Chapman & Hall, London, 1996.
- [26] Liu J, Wong W, Kong A. Covariance structure of the Gibbs sampler with applications to the comparisons of estimators and augmentation schemes. *Biometrika* 1994;81:27–40.
- [27] Schmid VJ, Whitcher B, Yang GZ, Taylor NJ, Padhani AR. Statistical analysis of pharmacokinetic models in dynamic contrast-enhanced magnetic resonance imaging. In J Duncan, G Gerig, eds.,

- Medical Imaging Computing and Computer-Assisted Intervention – MICCAI 2005, number 3750 in Lecture Notes in Computer Science. Berlin: Springer, 2005; 886–893.
- [28] Silverman BW. Density Estimation. London: Chapman & Hall, 1986.
- [29] Rue H. Fast sampling of Gaussian Markov random fields. Journal of the Royal Statistical Society Series B 2001;63:325–338.

Figure Captions

1. Schematic overview of the Bayesian hierarchical model for the observed contrast agent concentration time curves.
2. Contrast concentration time curves (CTCs) for pre- and post-treatment scans in three different subjects. Solid lines indicate the voxel-specific fit to the CTC, dashed lines the combined patient- and treatment-specific CTCs, and dotted lines the global pre- and post-treatment median CTCs for the entire study.
3. Samples from the marginal posterior distributions of K^{trans} at the study level. At pre-treatment K^{trans} is given by $\exp(\alpha_1)$ and at post-treatment K^{trans} is given by $\exp(\alpha_1 + \beta_1)$.
4. Samples from the marginal posterior distributions of K^{trans} at the patient level. At pre-treatment K^{trans} is given by $\exp(\alpha_1 + \gamma_{j1})$ for patient j and at post-treatment K^{trans} is given by $\exp(\alpha_1 + \beta_1 + \gamma_{j1} + \delta_{j1})$ for patient j .
5. Smoothed histograms summarizing the values of the posterior median K^{trans} at the voxel level. At pre-treatment K^{trans} is given by $\exp(\alpha_1 + \gamma_{j1} + \epsilon_{1jk1})$ for scan 1, patient j and voxel k . At post-treatment K^{trans} is given by $\exp(\alpha_1 + \beta_1 + \gamma_{1j} + \delta_{1j} + \epsilon_{12jk})$ for scan 2, patient j and voxel k . The x -axis has been restricted to $[0, 1]$ for visualization.
6. Smoothed histograms summarizing the values of K^{trans} from voxel-wise non-linear regression analysis. The x -axis has been restricted to $[0, 1]$ for visualization.

Table 1: Median K^{trans} values from the standard analysis (R = responder, NR = non-responder).

Patient ID	1	2	3	4	5	6	7	8	9	10	11	12
pathological	R	R	R	NR	NR	R	NR	NR	R	NR	NR	R
pre	0.208	0.355	0.255	0.230	0.199	0.154	0.264	0.198	0.305	0.267	0.432	0.174
post	0.161	0.120	0.031	0.245	0.208	0.173	0.327	0.223	0.122	0.221	0.111	0.113

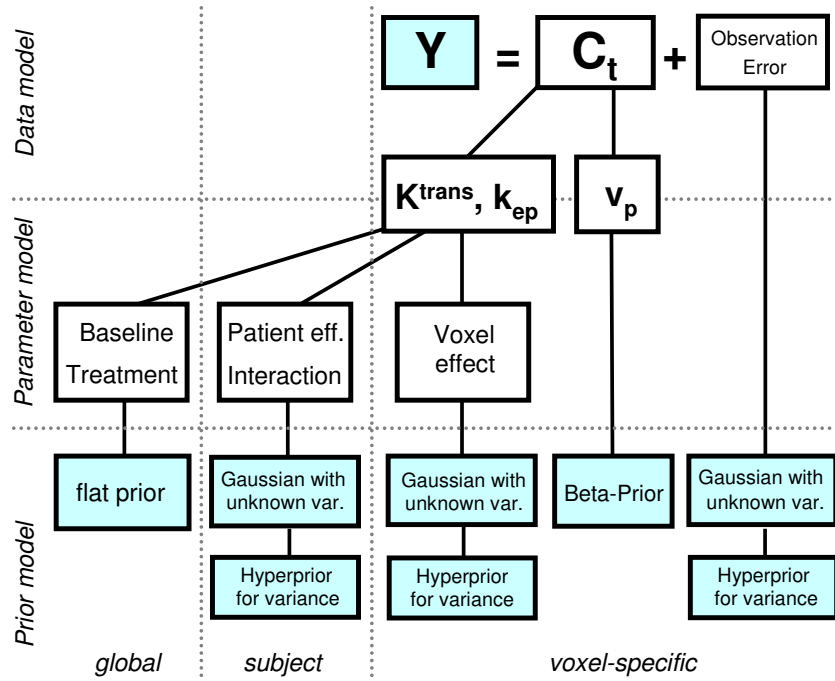


Figure 1:

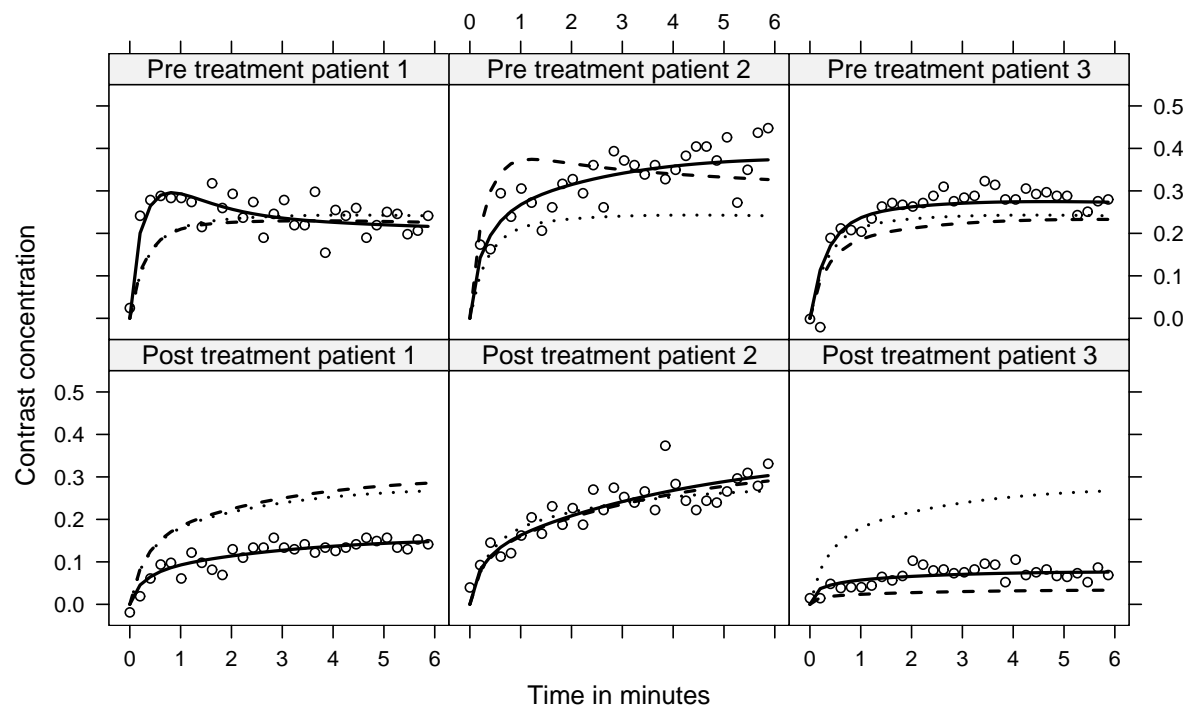


Figure 2:




Figure 3:




Figure 4:




Figure 5:




Figure 6: